CMS Tier 3 Organization & Plans

Rob Snihur
University of Nebraska, Lincoln

OSG All Hands Meeting: Joint Tier 3 Session
Fermilab

CMS Tier 3 Workshop

- Thanks to Ruth, Jemise, and the organizing committee
- Two sessions this morning. Discussions led by:
  - Malina Kirn (Maryland): Use of a Tier 3
  - Doug Johnson (Colorado): Set up of a Tier 3
  - Kevin L. Buterbaugh (Vanderbilt): Storage
  - Bill Strossman (Riverside): Hardware

USCMS Tier 3 Overview

- About 30 US Tier 3 sites exist
  - Various hardware/software configurations & support levels
  - Expect many more in the next year
  - GRID-enabled with priority to local users
- Goals:
  - Easy startup & monitoring
  - Minimize admin while operating (0.25 FTE)
  - Efficient data analysis (& MC production)

List of US Tier 3 Sites

Baylor	Brown	Cal Tech
Colorado	Cornell	FIT
FIU	FSU	Johns Hopkins
Iowa	Kansas	Kansas State
Maryland	Minnesota	Mississippi
Notre Dame	Ohio State	Omaha
Princeton	Pudue-Calumet	Rice
Riverside	Rutgers	Texas Tech
UCLA	UC Davis	Tennessee
UCSB	Vanderbilt	Virginia
Wayne State		

Tier 3 Use Cases

Malina Kirn (Maryland)

- Get data using CMS tool PhEDEx
- Analysis
  - Full CMS framework: cmsRun exe, submit via CMS Remote Analysis Builder (CRAB)
  - ROOT/PAT ntuple analyses
- Monte Carlo production
  - opportunistic
- Derive alignment, calibration constants
  - Short intense projects

Support & Documentation

- Community support model
- Biweekly meetings on EVO
- CMS Hyper News: https://hypernews.cern.ch/HyperNews/CMS/get/osg-tier3.html
- CMS Twiki: https://twiki.cern.ch/twiki/bin/viewauth/CMS/USTier3Computing
- Malina's instructions: http://hep-t3.physics.umd.edu/HowToForAdmins/index.html
- OSG Twiki (Marco Mambelli): https://twiki.grid.iu.edu/bin/view/Tier3/WebHome
- OSG Campfire
- GOC & Savannah tickets

Tier 3 setup

Doug Johnson (Colorado)

What are the requirements to be a T3?

Number of compute nodes

Sites range from a few to 100s of cores

Disk space

Lots of variation here are well

Network connection

Probably want Gb for data import, plenty of T3s with 100Mb

General thought is there are no requirements.

Components of a T3 – What do really need?

Batch system, Compute element, Storage Element, PhEDEx, GUMS/grid map file, SQUID/Frontier server, CMSSW

Storage Options

Kevin Buterbaugh (Vanderbilt)

- Simple NFS sharing to cluster
- Distributed filesystems:
  - HADOOP (see Brian Bockelman's talk)
  - Xrootd
  - REDDnet
- Dedicated storage systems:
  - Lustre
  - GPFS
- Pros & cons for each
  - Capacity, speed & efficiency, simplicity & ease

CMS Data Transfer Tool: PhEDEx

- Minimum for a Tier 3 (by some definitions)
- Rates vary between a few to 100+ MB/s
  - TCP tuning (Esnet)
  - Memory use lcg-cp vs. srm-copy
- Can be difficult to install & configure
- Planning a central PhEDEx service (Doug Johnson, Burt Holzman, RS)
  - Use by a Tier-3 site on a voluntary basis
  - FTS instead of SRM

Test bed

- Mini Tier-3 cluster at FNAL
- Virtual machines in xen & kvm
- CE, SEs (Tanya Levshina), PhEDEx, etc
- Allows rapid testing of configurations
- Performance mixed
- Distribute VMs for some services?

Future Plans

- CRAB developments
  - GlideIn & glite submission via CrabServer
- GUMS instead of gridmap
- Central PhEDEx server
- PhEDEx workshop
- Storage: HADOOP/REDDnet/LUSTRE/xrootd
- PROOF? analyze local data

Backup

Tier 3 Facility questions

Q: How many cores; storage; service nodes; % to be used by CMS; network speed; support personnel; GRID use; OSG stack?

General A: 10-100 cores; 10-100 TB; few service nodes; mainly CMS use; 1 Gbps; <~0.5 FTE; GRID enabled; use OSG stack

# cores	TB	# service Nodes	% for CMS	Gbps	# FTE	GRID/OSG	
30→192	4→48	1→2	100% CMS	1	1?	Y; Y	Brown
96	70	?	100%	20→100	use T2	N; N	CalTech
166→326	100-150	7-10	HEP-ex	1	0.25	Y; Y	Colorado
60	10	6	CMS priority	1→10	1 faculty	Y; Y	FIU
120	20-24	?	50%+ CMS	10	1?	Y; Y	FSU
8→32	20	use T2	100% CMS	use T2	use T2	Y; Y	KSU
120 nodes	15	?	CMS priority	10	1-2 stdnts	Y; Y	Iowa
64→128	10	2	100% CMS	?	0.5	Y; Y	Maryland
170	100	3	CMS priority	1	0.2+	Y; Y	Minnesota
80	100	use T2	use T2	?	use T2	Y; Y	MIT
113	25	1	90% CMS	1	campus	N; Y	Princeton
16	?	4-6	shared	?	~0	Y; Y	Rice
40→120	50	3	CMS priority	10	0.3	Y; Y	Riverside
80	50	?	100% CMS	0.4	1	Y; Y	Rutgers
8	?	?	100% CMS	?	?	Y; Y	TAMU
240+3360	45	8	shared by VO	1→10	0.5	Y; Y	TTU
128	20	?	100% CMS	?	limited	?; ?	UCSB
1 node	3	1	CMS only	1	0.1	N; N	Wayne State